

ParaMEDIC Enables Worldwide Supercomputer for Bioinformatics

Newswise — ParaMEDIC, a general software-based framework for large-scale distributed computing developed by Argonne National Laboratory (ANL) and Virginia Tech, will have a significant impact on the study of genomics.

The GenBank database, a collection of all publicly available DNA sequences, doubles in size approximately every 12 months. Yet the computational capability of a compute node doubles only every 18 months. Consequently, searching for similarities between new protein or nucleotide sequences with the database of known sequences is becoming increasingly difficult.

This problem becomes even more insidious when attempting a large-scale sequence-searching problem, like identifying missing genes, because solving such a problem oftentimes exceeds the computational and storage resources of any given supercomputing site. Thus, a team of researchers from ANL and Virginia Tech created a worldwide supercomputer called CompuMatrix and developed a novel framework called ParaMEDIC, short for Parallel Metadata Environment for Distributed I/O and Computing, in order to accelerate the speed of parallelized bioinformatics programs by an additional 27-fold when running over CompuMatrix.

ParaMEDIC achieves this performance by decoupling computation and storage in CompuMatrix. Pavan Balaji, post-doctoral researcher in the mathematics and computer science (MCS) division at ANL, and Wu-chun Feng, associate professor in the departments of computer science (CS) and electrical and computer engineering (ECE) at Virginia Tech, led the team of researchers who created ParaMEDIC.

With ParaMEDIC, a team of researchers led by Balaji and Feng embarked on two compute-intensive and storage-intensive tasks: sequence-searching all the known microbial genomes against each other in order to discover missing genes via mpiBLAST sequence-similarity computations (i.e., mpiBLAST, <http://www.mpiblast.org/> at Virginia Tech); and generating a complete genome similarity tree, based on the results of sequence-searching the above microbial genomes, in order to speed-up future sequence searches.

These two tasks required more than 10,000 processors across six supercomputing centers in the U.S. and generated a petabyte of uncompressed data in one month, which was then written in compressed format to a 0.5-petabyte filesystem in Japan.

With respect to discovering missing genes, João Setubal, associate professor and deputy director of the Virginia Bioinformatics Institute (VBI) at Virginia Tech, notes that most of the genomes completed to date have had their genes detected by gene-finder programs, which may miss real genes. "One way to discover these missed genes is by similarity computations," said Setubal. "If enough computer power is available, every possible location along a genome

can be checked for the presence of genes. That is exactly what the ParaMEDIC team has done by leveraging the mpiBLAST sequence-search program.”

In addition to discovering missing genes, the ParaMEDIC team also sought to restructure the microbial genome database by generating a complete genome sequence-similarity tree, thus enabling future searches to completion in just a fraction of time. “The sequence similarity tree can allow researchers to come up with better ways of structuring the database and more efficient algorithms to search specific portions of the data, instead of a brute-force search across the entire database as is currently done,” said Balaji. “Biologists can quickly discard huge parts of the database without losing any useful information.”

“Unfortunately, generating these similarity trees require large amounts of compute power and fast disk storage. No single supercomputing center could provide both the needed computational and storage resources needed to complete the above two tasks,” noted Feng. Thus, computations would have to be performed across a multitude of supercomputing sites and then petabytes of generated data (i.e., over 213,000 DVDs worth of storage) would need to be moved from the computational sites to another site for storage. “Such a model is clearly inefficient,” Feng said.

ParaMEDIC solves this problem by converting the output generated to orders-of-magnitude smaller metadata at the computation sites, transferring the metadata to the storage site, and then converting the metadata back to the actual output at the storage site. The metadata corresponding to the output will be available at Argonne National Laboratory and Virginia Tech for free public download in the future.

“Simultaneously needing large compute and storage resources has been a big problem for many applications,” said Ewing Lusk, director of the MCS Division at Argonne National Laboratory. “ParaMEDIC has provided a way to store relatively small metadata that can be used to guide future computations and regenerate the required portions of the actual output on-the-fly with very little computation. With the rapidly growing scales and sizes of applications, this model will be the wave of the future in a variety of compute- and storage-intensive applications,” said Lusk.

Six institutions provided computational resources for the CompuMatrix worldwide supercomputer: Virginia Tech, ANL, the Center for Computation and Technology at Louisiana State University, the Renaissance Computing Institute, the University of Chicago, and the San Diego Supercomputing Center.

The Tokyo Institute of Technology with support from Sun Microsystems generously provided the massive storage resources and I/O compute servers for the CompuMatrix worldwide supercomputer.

The ParaMEDIC framework is also one of the finalists in the Storage Challenge competition at this year’s IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC/07).